



CloudLab





The Need Addressed by CloudLab

- Clouds are changing the way we look at a lot of problems
 - Impacts go far beyond Computer Science
- ... but there's still a lot we don't know, from perspective of
 - Researchers (those who will transform the cloud)
 - Users (those who will use the cloud to transform their own fields)
- To investigate these questions, we need:
 - Flexible, scalable **scientific infrastructure**
 - That enables exploration of **fundamental** science in the cloud
 - Built **by** and **for** the research community



The CloudLab Vision

- A “meta-cloud” for building clouds
- Build your own cloud on our hardware resources
- Agnostic to specific cloud software
 - Run existing cloud software stacks (like OpenStack, Hadoop, etc.)
 - ... or new ones built from the ground up
- Control and visibility all the way to the bare metal
- “Sliceable” for multiple, isolated experiments at once

With CloudLab, it will be as easy to get a cloud tomorrow as it is to get a VM today



What Is CloudLab?

Slice A

*Geo-Distributed
Storage Research*

Slice B

*Stock
OpenStack*

- Supports transformative cloud research
- Built on Emulab and GENI
- Control to the bare metal
- Diverse, distributed resources
- Repeatable and scientific

Slice C

*Virtualization and
Isolation Research*

Slice D

*Allocation and Scheduling Research
for Cyber-Physical Systems*

Utah

Wisconsin

Clemson

GENI

CC-NIE, Internet2 AL2S, Regionals



Cloud Architecture Research

- Exploring **emerging and extreme** cloud architectures
- Evaluating **design choices** that exercise hardware and software capabilities
- Studying **geo-distributed** data centers for low-latency applications
- Developing different **isolation** models among tenants
- Quantifying **resilience** properties of architectures
- Developing new **diagnostic** frameworks
- Exploring cloud architectures for **cyber-physical systems**
- Enabling **realtime** and near-realtime compute services
- Enabling data-intensive computing (“**big data**”) at high performance in the cloud



Application Research Questions

- Experiment with **resource allocation** and scheduling
- Develop enhancements to **big data frameworks**
- Intra- and inter-datacenter **traffic engineering** and routing
- New tenant-facing **abstractions**
- New **mechanisms** in support of cloud-based services
- Study adapting **next-generation stacks** to clouds
- New troubleshooting and **anomaly detection** frameworks
- Explore different degrees of **security** and isolation
- **Composing** services from heterogeneous clouds
- **Application-driven** cloud architectures



CloudLab's Hardware

One facility, one account, three locations

- About 5,000 cores each (15,000 total)
- 8-16 cores per node
- Baseline: 4GB RAM / core
- Latest virtualization hardware
- TOR / Core switching design
- 10 Gb to nodes, SDN
- 100 Gb to Internet2 AL2S
- *Partnerships with multiple vendors*

Wisconsin

- **Storage and net.**
- Per node:
 - 128 GB RAM
 - 2x1TB Disk
 - 400 GB SSD
- Clos topology
- *Cisco*

Clemson

- **High-memory**
- 16 GB RAM / core
- 16 cores / node
- Bulk block store
- Net. up to 40Gb
- High capacity
- *Dell*


Utah

- **Power-efficient**
- ARM64 / x86
- Power monitors
- Flash on ARM64s
- Disk on x86
- Very dense
- *HP*



Technology Foundations

- Built on Emulab and GENI (“ProtoGENI”)
- In active development at Utah since 1999
- Several thousand users (incl. GENI users)
- Provisions, then gets out of the way
 - “Run-time” services are optional
- Controllable through a web interface and GENI APIs
- *Scientific instrument for repeatable research*
 - Physical isolation for most resources
 - *Profiles* capture everything needed for experiments
 - Software, data, and hardware details
 - Can be shared and published (eg. in papers)



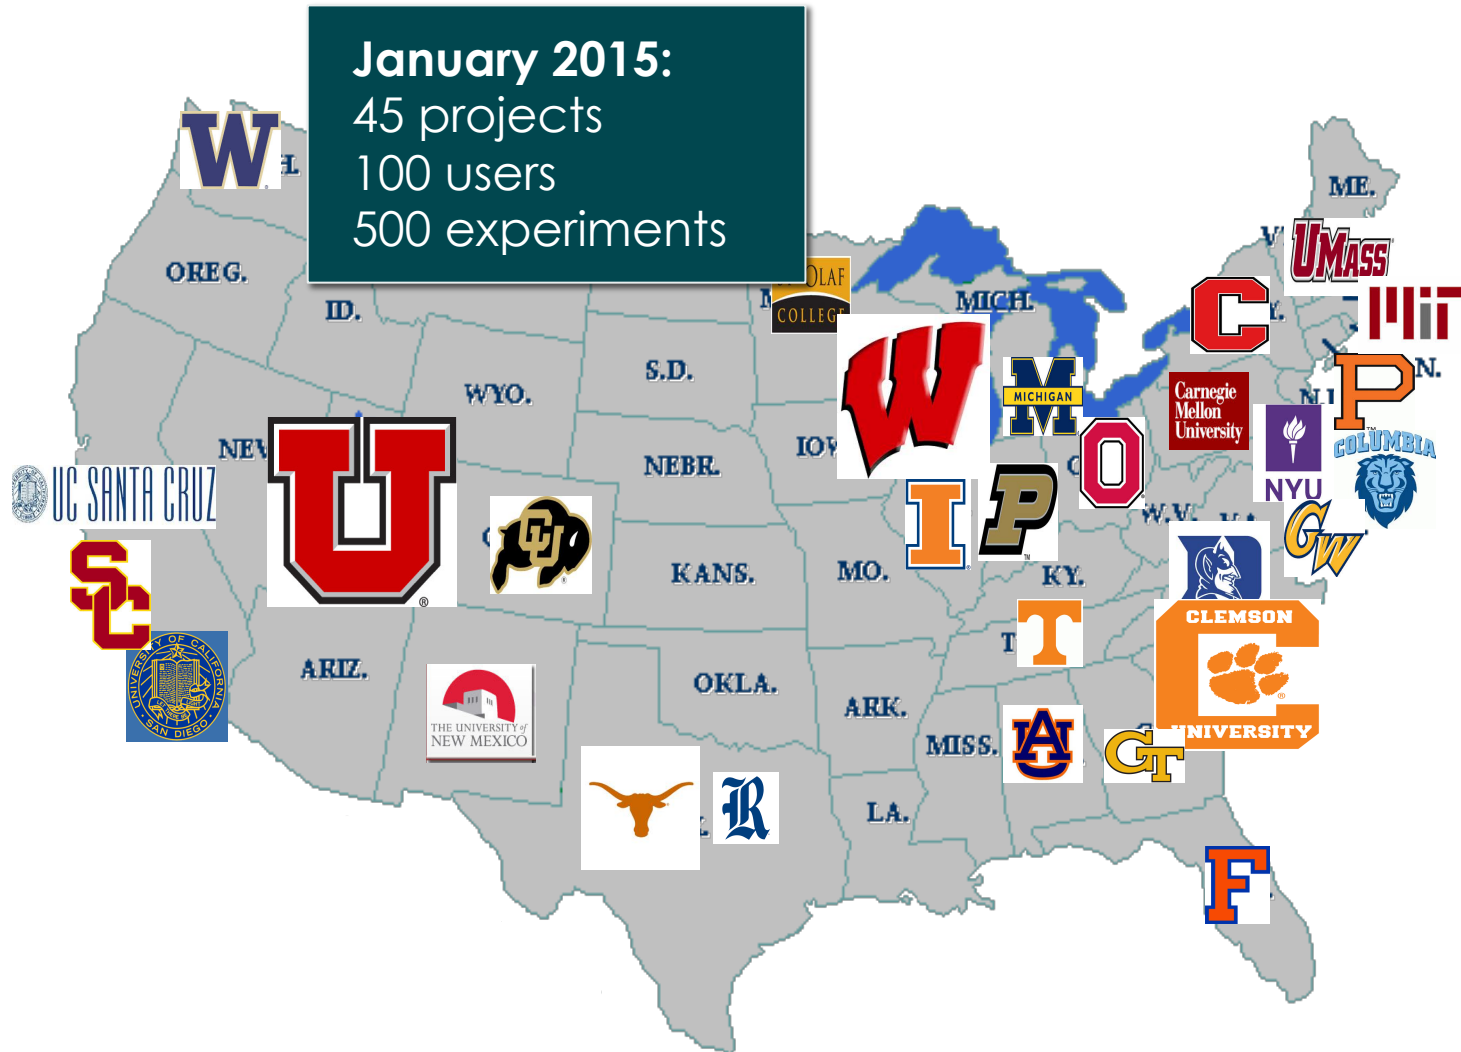


Who can use CloudLab?

- US academics and educators
 - Researchers in cloud architecture and novel cloud applications
 - Teaching classes, other training activities
- No charge: free for research and educational use
- International federations expected
- Apply on the website at www.cloudlab.us
 - Everyone with a GENI or Emulab account (several thousand)




January 2015:
45 projects
100 users
500 experiments





Federated with GENI

- *CloudLab can be used with a GENI account, and vice-versa*
- GENI Racks: ~ 50 small clusters around the country
- Programmable wide-area network
 - Openflow at dozens of sites
 - Connected in one layer 2 domain
- Large clusters (100s of nodes) at several sites
- Wireless and mobile
 - WiMax at 8 institutions
 - LTE / EPC testbed (“PhantomNet”) at Utah
- International partners
 - Europe (FIRE), Brazil, Japan





Community Outreach

- Applications in areas of national priority
 - Medicine, emergency response, smart grids, etc.
 - Through 
- “Opt in” to compute jobs from domain scientists



- Summer camps
 - Through Clemson data-intensive computing program
- Under-represented groups





Availability and Schedule

- Availability:
 - ✓ **Open and in use!**
- Hardware being deployed in stages:
 - ✓ Fall 2014: Utah / HP cluster
 - Winter 2015: Wisconsin / Cisco cluster
 - Winter 2015: Dell / Clemson cluster
- Hardware refreshes in 2015 and 2016



Your Own Cloud in One Click

The screenshot shows a web browser window with the URL `cloudlab.us`. The page has a dark teal header with navigation links: [Home](#), [Manual](#), and [Actions](#) (with a dropdown arrow). The CloudLab logo is centered in the header, and on the right, it says "rpruser logged in" next to a [Logout](#) button. The main content area has the heading "What is CloudLab?" followed by "Run an Experiment". A box displays the "Selected Profile: ARM64OpenStack" with a description: "An ARM64 OpenStack instance with a controller, network manager, and three compute nodes." Below this are "Copy Profile" and "Change Profile" buttons. A "Cluster:" label is next to a dropdown menu showing "Utah Cloudlab". A large green "Create!" button is at the bottom of the form. The footer contains "Powered by emulab", "Question or comment? Join the Help Forum", and "Supported by NSF © 2014 The University of Utah".

cloudlab.us

Home Manual Actions

rpruser logged in Logout

What is CloudLab?

Run an Experiment

Selected Profile: ARM64OpenStack

An ARM64 OpenStack instance with a controller, network manager, and three compute nodes.

Copy Profile Change Profile

Cluster: Utah Cloudlab

Create!

Powered by emulab Question or comment? Join the Help Forum Supported by NSF © 2014 The University of Utah



The CloudLab Team



Robert Ricci (PI)
Eric Eide
Steve Corbató
Kobus Van der Merwe



Aditya Akella (co-PI)
Remzi Arpaci-Dusseau
Miron Livny



KC Wang (co-PI)
Jim Bottum
Jim Pepin
Amy Apon



Chip Elliott (co-PI)
Larry Landweber



Mike Zink (co-PI)
David Irwin



Glenn Ricart (co-PI)





Challenges We Can Use Help With

- Looking down the road: where is cloud research going?
 - Applications, software layer, hardware, software, new use cases, service models, etc.
- Build profiles for cloud software stacks
 - Both general, and those used by particular communities
 - Standard platforms to build innovations on
- Study the facility itself
 - What data would be useful for us to collect and make available?
- Merging cloud and other computing paradigms
 - Our provisioning system is not limited to cloud



Learn more, sign up:

www.CloudLab.us



This material is based upon work supported by the National Science Foundation under Grant No. 1419199. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.